



Configuring the System

This chapter provides information about changing switch-wide configuration settings. It includes command-line interface (CLI) procedures for using commands that have been specifically created or changed for the Catalyst 2950 switches. For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 2950 Desktop Switch Command Reference*.

This chapter does not repeat the concepts and CLI procedures provided in the standard Cisco IOS Release 12.1 documentation. For information about the standard IOS Release 12.1 commands, refer to the IOS documentation set available from the Cisco.com home page at **Service and Support > Technical Documents**. On the Cisco Product Documentation home page, select Release 12.1 from the Cisco IOS Software drop-down list.

For information about configuring these settings from Cluster Management Suite (CMS), refer to the online help.



Note

Some features can be implemented only by using the CLI.

Changing IP Information

You can assign and change the IP information of your switch in these ways:

- Using the setup program, as described in the release notes
- Manually assigning an IP address, as described in this section
- Using Dynamic Host Configuration Protocol (DHCP)-based autoconfiguration, as described in this section



Caution

Changing the switch IP address ends any CMS, Telnet, or Simple Network Management Protocol (SNMP) session. To restart your CMS session, enter the new IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Internet Explorer). To restart your CLI session through Telnet, follow the steps described in the [“Accessing the CLI” section on page 3-9](#).



Note

If you enabled the DHCP feature, the switch assumes you are using an external server for IP address allocation. While this feature is enabled, any values you manually enter (from the CMS or from the **ip address** command) are ignored.

Manually Assigning and Removing Switch IP Information

You can manually assign an IP address, mask, and default gateway to the switch. The mask identifies the bits that denote the network number in the IP address. When you use the mask to subnet a network, the mask is then referred to as a subnet mask. The broadcast address is reserved for sending messages to all hosts. The CPU sends traffic to an unknown IP address through the default gateway.

Beginning in privileged EXEC mode, follow these steps to enter the IP information:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface vlan 1	Enter interface configuration mode, and enter the VLAN to which the IP information is assigned. VLAN 1 is the default management VLAN, but you can configure any VLAN from 1 to 1001.
Step 3	ip address <i>ip_address subnet_mask</i>	Enter the IP address and subnet mask.
Step 4	exit	Return to global configuration mode.
Step 5	ip default-gateway <i>ip_address</i>	Enter the IP address of the default router.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify that you entered the information was entered correctly by displaying the running configuration. If the information is incorrect, repeat the procedure.

Use this procedure to remove the IP information from a switch.



Note

Using the **no ip address** command in configuration mode disables the IP stack as well as removes the IP information. Cluster members without IP addresses rely on the enabled IP stack.

Beginning in privileged EXEC mode, follow these steps to remove an IP address:

	Command	Purpose
Step 1	no ip address <i>ip_address subnet_mask</i>	Remove the IP address and subnet mask.
Step 2	end	Return to privileged EXEC mode.
Step 3	show running-config	Verify that you entered the information was removed by displaying the running configuration.

Using DHCP-Based Autoconfiguration

The Dynamic Host Configuration Protocol (DHCP) provides configuration information to Internet hosts and internetworking devices. With DHCP-based autoconfiguration, your switch (DHCP client) can be automatically configured during bootup with IP address information and a configuration file that it receives during DHCP-based autoconfiguration.

**Note**

DHCP replaces the Bootstrap Protocol (BOOTP) feature autoconfiguration to ensure retrieval of configuration files by unicast TFTP messages. BOOTP is available in earlier software releases for this switch.

Understanding DHCP-Based Autoconfiguration

The DHCP provides configuration information to Internet hosts and internetworking devices. This protocol consists of two components: one for delivering configuration parameters from a DHCP server to a device and one for allocating network addresses to devices. DHCP is built on a client-server model, in which designated DHCP servers allocate network addresses and deliver configuration parameters to dynamically configured devices.

With DHCP-based autoconfiguration, your switch (DHCP client) can be automatically configured at startup with IP address information and a configuration file that it receives during DHCP-based autoconfiguration. No DHCP client-side configuration is required on your switch.

However, you need to configure the DHCP server for various lease options. You might also need to configure a TFTP server, a Domain Name System (DNS) server, and possibly a relay device if the servers are on a different LAN than your switch. A relay device forwards broadcast traffic between two directly connected LANs. A router does not forward broadcast packets, but it forwards packets based on the destination IP address in the received packet. DHCP-based autoconfiguration replaces the BOOTP client functionality on your switch.


DHCP Client Request Process

When you boot your switch, the DHCP client can be invoked and automatically request configuration information from a DHCP server under these conditions:

- The configuration file is not present on the switch.
- The configuration file is present, but the IP address is not specified in it.
- The configuration file is present, the IP address is not specified in it, and the **service config** global configuration command is included. This command enables the auto-loading of a configuration file from a network server.

Figure 6-1 shows the sequence of messages that are exchanged between the DHCP client and the DHCP server.

Figure 6-1 DHCP Request for IP Information from a DHCP Server



The client, Switch A, broadcasts a DHCPDISCOVER message to locate a DHCP server. The DHCP server offers configuration parameters (such as an IP address, subnet mask, gateway IP address, DNS IP address, a lease for the IP address, and so forth) to the client in a DHCPOFFER unicast message.

In a DHCPREQUEST broadcast message, the client returns a request for the offered configuration information to the DHCP server. The request is broadcast so that all other DHCP servers that received the DHCPDISCOVER broadcast message from the client can reclaim the IP addresses that they offered to the client.

The DHCP server confirms that the IP address has been allocated to the client by returning a DHCPACK unicast message to the client. With this message, the client and server are bound, and the client uses configuration information received from the server. The amount of information the switch receives depends on how you configure the DHCP server. For more information, see the “[Configuring the DHCP Server](#)” section on page 6-4.

If the configuration parameters sent to the client in the DHCP OFFER unicast message by the DHCP server are invalid (a configuration error exists), the client returns a DHCPDECLINE broadcast message to the DHCP server.

The DHCP server sends the client a DHCPNAK denial broadcast message, which means that the offered configuration parameters have not been assigned, that an error has occurred during the negotiation of the parameters, or that the client has been slow in responding to the DHCP OFFER message (the DHCP server assigned the parameters to another client) of the DHCP server.

A DHCP client might receive offers from multiple DHCP or BOOTP servers and can accept any of the offers; however, the client usually accepts the first offer it receives. The offer from the DHCP server is not a guarantee that the IP address will be allocated to the client; however, the server usually reserves the address until the client has had a chance to formally request the address. If the switch accepts replies from a BOOTP server and configures itself, the switch will broadcast, instead of unicast, TFTP requests to obtain the switch configuration file.

Configuring the DHCP Server

You should configure the DHCP servers with reserved leases that are bound to each switch by the switch hardware address. If the DHCP server does not support reserved leases, the switch can obtain different IP addresses and configuration files at different boot instances. You should configure the DHCP server with these lease options:

- IP address of the client (required)
- Subnet mask of the client (required)
- DNS server IP address (required)
- Router IP address (default gateway address to be used by the switch) (required)
- TFTP server name (required)
- Boot filename (the name of the configuration file that the client needs) (recommended)
- Host name (optional)

If you do not configure the DHCP server with the lease options described earlier, it replies to client requests with only those parameters that have available values. If the IP address and subnet mask are not in the reply, the switch is not configured. If the DNS server IP address, router IP address, or TFTP server name are not found, the switch might broadcast TFTP requests. Unavailability of other lease options does not affect autoconfiguration.

**Note**

If the configuration file on the switch does not contain the IP address, the switch obtains its address, mask, gateway IP address, and host name from DHCP. If the **service config** global configuration command is specified in the configuration file, the switch receives the configuration file through TFTP requests. If both the **service config** global configuration command and the IP address are in the configuration file, DHCP is not used, and the switch obtains the default configuration file by broadcasting TFTP requests.

The DHCP server can be on the same or a different LAN as the switch. If it is on a different LAN, the switch must be able to access it through a relay device.

For more information, see the [“Configuring the Relay Device” section on page 6-6](#). You must also set the TFTP server with the switch configuration files; for more information, see the next section.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Configuring the TFTP Server

The TFTP server must contain one or more configuration files in its base directory. The files can include these:

- The configuration file named in the DHCP reply (the actual switch configuration file).
- The network-config or the cisco.net.cfg file (known as the default configuration files).
- The router-config or the ciscortr.cfg file (These files contain commands common to all switches. Normally, if the DHCP and TFTP servers are properly configured, these files are not accessed.)

You must specify the TFTP server name in the DHCP-server lease database. You must also specify the TFTP server name-to-IP-address mapping in the DNS-server database.

The TFTP server can be on the same or on a different LAN as the switch. If it is on a different LAN, the switch must be able to access it through a relay device or a router. For more information, see the [“Configuring the Relay Device” section on page 6-6](#).

If the configuration filename is provided in the DHCP server reply, the configuration files for a switch can be spread over multiple TFTP servers. However, if the configuration filename is not provided, the configuration files must reside on a single TFTP server.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Configuring the Domain Name and the DNS

Each unique IP address can have a host name associated with it. The IOS software maintains a cache of host name-to-address mappings for use by the EXEC mode **connect**, **telnet**, and **ping** commands, and related Telnet-support operations. This cache speeds the process of converting names to addresses.

IP defines a hierarchical naming scheme that allows a device to be identified by its location or domain. Domain names use periods (.) as the delimiting characters. For example, Cisco Systems is a commercial organization that IP identifies by a *com* domain name, so its domain name is *cisco.com*. A specific device in this domain, the File Transfer Protocol (FTP) system for example, is identified as *ftp.cisco.com*.

To keep track of domain names, IP has defined the concept of a Domain Name Server (DNS), which holds a cache (or database) of names mapped to IP addresses. To map domain names to IP addresses, you must first identify the host names and then specify a name server and enable the DNS, the Internet’s global naming scheme that uniquely identifies network devices.

You can specify a default domain name that the software uses to complete domain name requests. You can specify either a single domain name or a list of domain names. When you specify a domain name, any IP host name without a domain name will have that domain name appended to it before being added to the host table.

If your network devices require connectivity with devices in networks for which you do not control name assignment, you can assign device names that uniquely identify your devices within the entire internetwork. The DNS accomplishes this task. This service is enabled by default.

The switch uses the DNS server to resolve the TFTP server name to a TFTP-server IP address. You must configure the TFTP server name-to-IP address map on the DNS server. The TFTP server contains the configuration files for the switch.

You must configure the IP addresses of the DNS servers in the lease database of the DHCP server from where the DHCP replies will retrieve them. You can enter up to two DNS server IP addresses in the lease database.

The DNS server can be on the same or on a different LAN as the switch. If it is on a different LAN, the switch must be able to access it through a relay device or router. For more information, see the [“Configuring the Relay Device”](#) section on page 6-6.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Configuring the Relay Device

You need to use a relay device if the DHCP, DNS, or TFTP servers are on a different LAN than the switch. You must configure this relay device to forward received broadcast packets on an interface to the destination host. This configuration ensures that broadcasts from the DHCP client can reach the DHCP, DNS, and TFTP servers and that broadcasts from the servers can reach the DHCP client.

If the relay device is a Cisco router, you enable IP routing (**ip routing** global configuration command) and configure it with helper addresses by using the **ip helper-address** interface configuration command.


For example, in [Figure 6-2](#), you configure the router interfaces as follows:

On interface 10.0.0.2:

```
router(config-if)# ip helper-address 20.0.0.2
router(config-if)# ip helper-address 20.0.0.3
router(config-if)# ip helper-address 20.0.0.4
```

On interface 20.0.0.1

```
router(config-if)# ip helper-address 10.0.0.1
```

Figure 6-2 Relay Device Used in Autoconfiguration

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Obtaining Configuration Files

Depending on the availability of the IP address and the configuration filename in the DHCP reserved lease, the switch obtains its configuration information in these ways:

- The IP address and the configuration filename is reserved for the switch and provided in the DHCP reply (one-file read method).

The switch receives its IP address, subnet mask, and configuration filename from the DHCP server. It also receives a DNS server IP address and a TFTP server name. The switch sends a DNS request to the DNS server, specifying the TFTP server name, to obtain the TFTP server address. Then the switch sends a unicast message to the TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, completes its boot-up process.

- Only the configuration filename is reserved for the switch. The IP address is dynamically allocated to the switch by the DHCP server (one-file read method).

The switch follows the same configuration process described in the first item.

- Only the IP address is reserved for the switch and provided in the DHCP reply. The configuration filename is not provided (two-file read method).

The switch receives its IP address and subnet mask from the DHCP server. It also receives a DNS server IP address and a TFTP server name. The switch sends a DNS request to the DNS server, specifying the TFTP server name, to obtain the TFTP server address.

The switch sends a unicast message to the TFTP server to retrieve the network-config or cisco.net.cfg default configuration file. (If the network-config file cannot be read, the switch reads the cisco.net.cfg file.)

The default configuration file contains the host names-to-IP-address mapping for the switch. The switch fills its host table with the information in the file and obtains its host name. If the host name is not found in the file, the switch uses the host name in the DHCP reply. If the host name is not specified in the DHCP reply, the switch uses the default “Switch” as its host name.

After obtaining its host name from the default configuration file or the DHCP reply, the switch reads the configuration file that has the same name as its host name (*hostname-conf* or *hostname.cfg*, depending on whether *network-conf* or *cisconet.cfg* was read earlier) from the TFTP server. If the *cisconet.cfg* file is read, the filename of the host is truncated to eight characters.

If the switch cannot read the *network-conf*, *cisconet.cfg*, or the host-name file, it reads the *router-conf* file. If the switch cannot read the *router-conf* file, it reads the *ciscotr.cfg* file.

**Note**

The switch broadcasts TFTP server requests if the TFTP server name is not obtained from the DHCP replies, if all attempts to read the configuration file through unicast transmissions fail, or if the TFTP server name cannot be resolved to an IP address.

Example Configuration

Figure 6-3 shows a sample network for retrieving IP information by using DHCP-based autoconfiguration.

Figure 6-3 DHCP-Based Autoconfiguration Network Example




Table 6-1 shows the configuration of the reserved leases on the DHCP server.

Table 6-1 DHCP Server Configuration

	Switch-1	Switch-2	Switch-3	Switch-4
Binding key (hardware address)	00e0.9f1e.2001	00e0.9f1e.2002	00e0.9f1e.2003	00e0.9f1e.2004
IP address	10.0.0.21	10.0.0.22	10.0.0.23	10.0.0.24
Subnet mask	255.255.255.0	255.255.255.0	255.255.255.0	255.255.255.0
Router address	10.0.0.10	10.0.0.10	10.0.0.10	10.0.0.10
DNS server address	10.0.0.2	10.0.0.2	10.0.0.2	10.0.0.2
TFTP server name	<i>maritsu</i> or <i>10.0.0.3</i>	<i>maritsu</i> or <i>10.0.0.3</i>	<i>maritsu</i> or <i>10.0.0.3</i>	<i>maritsu</i> or <i>10.0.0.3</i>
Boot filename (configuration file) (optional)	switch1-confg	switch2-confg	switch3-confg	switch4-confg
Host name (optional)	switch1	switch2	switch3	switch4

DNS Server Configuration

The DNS server maps the TFTP server name *maritsu* to IP address 10.0.0.3.

TFTP Server Configuration (on UNIX)

The TFTP server base directory is set to `/tftpserver/work/`. This directory contains the `network-confg` file used in the two-file read method. This file contains the host name to be assigned to the switch based on its IP address. The base directory also contains a configuration file for each switch (`switch1-confg`, `switch2-confg`, and so forth) as shown in this display:

```
prompt> cd /tftpserver/work/
prompt> ls
network-confg
switch1-confg
switch2-confg
switch3-confg
switch4-confg
prompt> cat network-confg
ip host switch1 10.0.0.21
ip host switch2 10.0.0.22
ip host switch3 10.0.0.23
ip host switch4 10.0.0.24
```

DHCP Client Configuration

No configuration file is present on Switch 1 through Switch 4.

Configuration Explanation

In [Figure 6-3](#), Switch 1 reads its configuration file as follows:

- It obtains its IP address 10.0.0.21 from the DHCP server.
- If no configuration filename is given in the DHCP server reply, Switch 1 reads the `network-confg` file from the base directory of the TFTP server.
- It adds the contents of the `network-confg` file to its host table.
- It reads its host table by indexing its IP address 10.0.0.21 to its host name (switch1).

- It reads the configuration file that corresponds to its host name; for example, it reads switch1-config from the TFTP server.

Switches 2 through 4 retrieve their configuration files and IP addresses in the same way.

Changing the Password

You can assign the password of your switch in these ways:

- Using the setup program, as described in the release notes
- Manually assigning a password, as described in this section



Note

You can change a password only by using the CLI. Your connection with the switch ends when you change the enable secret password. You will then need to reopen the session with the new password. If you have forgotten your password, see the [“Recovering from a Lost or Forgotten Password” section on page 14-9](#).

Because many privileged EXEC commands are used to set operating parameters, you should password-protect these commands to prevent unauthorized use. Catalyst 2950 switches have two commands for setting passwords:

- **enable secret** *password* (a very secure, encrypted password)
- **enable password** *password* (a less secure, unencrypted password)

You must enter one of these passwords to gain access to privileged EXEC mode. We recommend that you use the enable secret password.



Note

When set, the enable secret password takes precedence, and the enable password serves no purpose.

If you enter the **enable secret** command, the text is encrypted before it is written to the config.text file, and it is unreadable. If you enter the **enable password** command, the text is written as entered to the config.text file where you can read it.

You can also specify up to 15 privilege levels and define passwords for them by using the **enable password [level level] {password}** or the **enable secret [level level] {password}** command. Level 1 is EXEC-mode user privileges. If you do not specify a level, the privilege level defaults to 15 (privileged EXEC-mode privileges).

You can specify a level, set a password, and give the password only to users who need to have access at this level. Use the **privilege level** global configuration command to specify commands accessible at various levels.



Note

You need an enable secret password with a privilege level 15 to access CMS. You must also use this password if you configure the Terminal Access Controller Access Control System Plus (TACACS+) protocol from the CLI so that all your HTTP connections are authenticated through the TACACS+ server. The Telnet password must be an enable secret password.

For information about managing passwords in switch clusters, see the [“Passwords” section on page 5-14](#).

Both types of passwords can contain from 1 to 25 uppercase and lowercase alphanumeric characters, and both can start with a number. Spaces are also valid password characters; for example, *two words* is a valid password. Leading spaces are ignored; trailing spaces are recognized. The password is case sensitive.

To remove a password, use the **no** version of the commands: **no enable secret** or **no enable password**. If you lose or forget your enable password, see the [“Recovering from a Lost or Forgotten Password” section on page 14-9](#).

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Setting the System Date and Time

You can change the date and a 24-hour clock time setting on the switch. If you are entering the time for an American time zone, enter the three-letter abbreviation for the time zone, such as PST for Pacific standard time. If you are identifying the time zone by referring to Greenwich mean time, enter UTC (universal coordinated time). You then must enter a negative or positive number as an offset to indicate the number of time zones between the switch and Greenwich, England. Enter a negative number if the switch is west of Greenwich, England, and east of the international date line. For example, California is seven time zones west of Greenwich, so you would enter -7 . Enter a positive number if the switch is east of Greenwich. You can also enter negative and positive numbers for minutes.

Configuring Daylight Saving Time

You can configure the switch to change to daylight saving time on a particular day every year, on a day that you enter, or not at all.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Configuring the Network Time Protocol

In complex networks, it is often prudent to distribute time information from a central server. The Network Time Protocol (NTP) can distribute time information by responding to requests from clients or by broadcasting time information.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Configuring the Switch as an NTP Client

You configure the switch as an NTP client by entering the IP addresses of up to ten NTP servers and specifying which server should be used first. You can also enter an authentication key to be used as a password when requests for time information are sent to the server.

Enabling NTP Authentication

To ensure the validity of information received from NTP servers, you can authenticate NTP messages with public-key encryption. This procedure must be coordinated with the administrator of the NTP servers: the information you enter will be matched by the servers to authenticate it.

Configuring the Switch for NTP Broadcast-Client Mode

You can configure the switch to receive NTP broadcast messages if there is an NTP broadcast server, such as a router, broadcasting time information on the network. You can also enter a value to account for any round-trip delay between the client and the NTP broadcast server.

Configuring SNMP

If your switch is part of a cluster, the clustering software can change Simple Network Management Protocol (SNMP) parameters (such as host names) when the cluster is created. If you are configuring a cluster for SNMP, see the [“SNMP Community Strings” section on page 5-14](#).

Disabling and Enabling SNMP

SNMP is enabled by default and must be enabled for Cluster Management features to work properly.

SNMP is always enabled for Catalyst 1900 and Catalyst 2820 switches.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Entering Community Strings

Community strings serve as passwords for SNMP messages, permitting access to the agent on the switch. If you are entering community strings for a cluster member, see the [“SNMP Community Strings” section on page 5-14](#). You can enter community strings with these characteristics:

Read-only (RO)—Requests accompanied by the string can display MIB-object information.

Read-write (RW)—Requests accompanied by the string can display MIB-object information and set MIB objects.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Adding Trap Managers

A trap manager is a management station that receives and processes traps. When you configure a trap manager, the community strings for each member switch must be unique. If a member switch has an assigned IP address, the management station accesses the switch by using that IP address.

By default, no trap manager is defined, and no traps are issued. [Table 6-2](#) describes the Catalyst 2950 switch traps. You can enable any or all of these traps and configure a trap manager on these switches to receive them.

Table 6-2 Catalyst 2950 Switch Traps

Config	Generate traps whenever the switch configuration changes.
SNMP	Generate the supported SNMP traps.
TTY	Generate traps when the switch starts a management console CLI session.
VLAN membership	Generate a trap for each VLAN Membership Policy Server (VMPS) change.
VTP	Generate a trap for each VLAN Trunking Protocol (VTP) change.
cluster	Generate the cluster traps.
entity	Generate the ENTITY_MIB traps.
hsrp	Generate the SNMP HSRP traps.
rtr	Enable the SNMP Response Time Reporter Traps.
mac-notification	Generate a trap when a MAC address is added or removed from any interface.
C2900/3500	Generate the switch-specific traps. These traps are in the private enterprise-specific Management Information Base (MIB).

Beginning in privileged EXEC mode, follow these steps to add a trap manager and a community string:

	Command	Purpose
Step 1	config terminal	Enter global configuration mode.
Step 2	snmp-server host 172.2.128.263 <i>community-string</i> snmp vlan-membership	Enter the trap manager IP address, the community string, and the traps to generate.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify that you entered the information correctly by displaying the running configuration.

Configuring CDP

Use the Cisco IOS CLI and Cisco Discovery Protocol (CDP) to enable CDP for the switch, set global CDP parameters, and display information about neighboring Cisco devices.

CDP enables the Cluster Management Suite (CMS) to display a graphical view of the network. For example, the switch uses CDP to find cluster candidates and to maintain information about cluster members and other devices up to three cluster-enabled devices away from the command switch.

You can configure CDP to discover switches running the CMS up to seven devices away from the command switch. Devices that do not run clustering software display as edge devices, and CDP cannot discover any device connected to them.



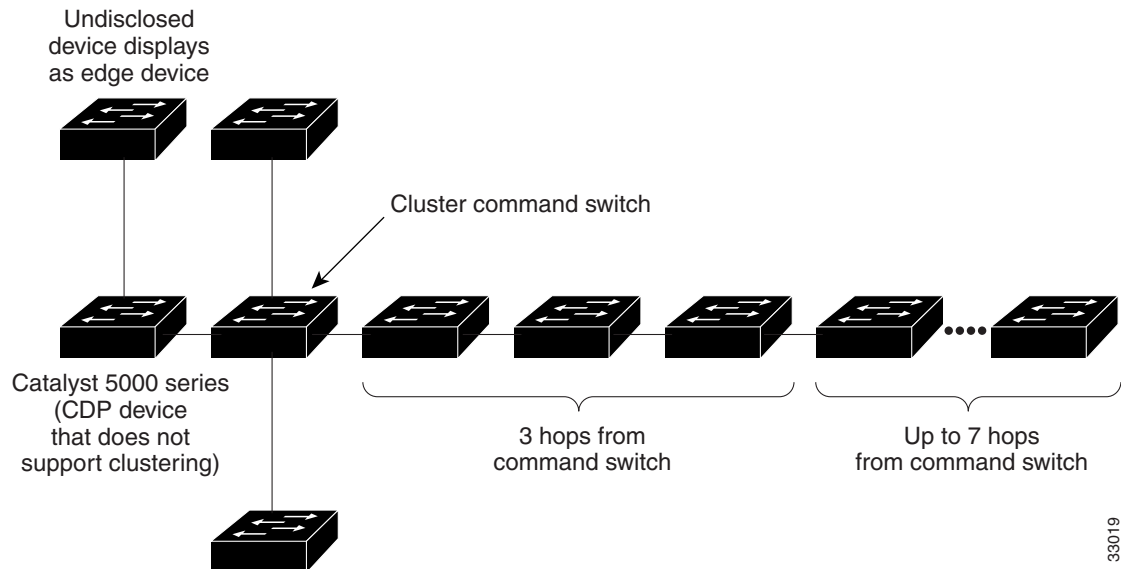
Note

Creating and maintaining switch clusters is based on the regular exchange of CDP messages. Disabling CDP can interrupt cluster discovery. For more information about the role that CDP plays in clustering, see the [“Automatic Discovery of Cluster Candidates and Members”](#) section on page 5-4.

Configuring CDP for Extended Discovery

You can change the default configuration of CDP on the command switch to discover devices up to seven *hops* away. See [Figure 6-4](#). [Figure 6-4](#) also shows the command switch connected to a Catalyst 5000 series switch. Although the Catalyst 5000 supports CDP, it does not support clustering, and the command switch cannot learn about connected candidate switches connected to it, even if they are running CMS.

Figure 6-4 Discovering Cluster Candidates through CDP



33019

Beginning in privileged EXEC mode, follow these steps to configure the number of hops that CDP uses to discover candidate switches and cluster members.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cluster discovery hop-count number	Enter the number of hops that you want CDP to search for cluster candidates and cluster members.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify the change by displaying the running configuration file. The hop count is displayed in the file.

Managing the ARP Table

To communicate with a device (over Ethernet, for example), the software first must determine the 48-bit MAC or the local data link address of that device. The process of determining the local data link address from an IP address is called *address resolution*.

The Address Resolution Protocol (ARP) associates a host IP address with the corresponding media or MAC addresses and the VLAN ID. Taking an IP address as input, ARP determines the associated MAC address. Once a MAC address is determined, the IP-MAC address association is stored in an ARP cache for rapid retrieval. Then the IP datagram is encapsulated in a link-layer frame and sent over the network. Encapsulation of IP datagrams and ARP requests and replies on IEEE 802 networks other than Ethernet is specified by the Subnetwork Access Protocol (SNAP). By default, standard Ethernet-style ARP encapsulation (represented by the **arpa** keyword) is enabled on the IP interface.

ARP entries added manually to the table do not age and must be manually removed.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com for additional information and CLI procedures.

Managing the MAC Address Tables

You can manage the MAC address tables that the switch uses to forward traffic between ports. All MAC addresses in the address tables are associated with one or more ports. These MAC tables include these types of addresses:

- **Dynamic address:** a source MAC address that the switch learns and then drops when it is not in use.
- **Secure address:** a manually entered unicast address or dynamically learnt address that is usually associated with a secured port. Secure addresses do not age.
- **Static address:** a manually entered unicast or multicast address that does not age and that is not lost when the switch resets.

The address tables list the destination MAC address and the VLAN ID, module, and port number associated with the address. This example shows the list of addresses as they would appear in the dynamic, secure, or static address table.

```

Mac Address Table
-----
Vlan      Mac Address      Type      Ports
-----
  1       0001.42e2.cdcd   DYNAMIC   Fa0/24
  1       0001.96e4.fed6   DYNAMIC   Fa0/2
  1       0030.19c6.54dd   DYNAMIC   Fa0/24
 10       0000.0000.0001   STATIC    Fa0/7
 10       0404.0400.0006   DYNAMIC   Fa0/7
Total Mac Addresses for this criterion:5

```

For information about the Mac address Notification feature, see the [“MAC Address Notification” section on page 6-17](#).

MAC Addresses and VLANs

All addresses are associated with a VLAN. An address can exist in more than one VLAN and have different destinations in each. Multicast addresses, for example, could be forwarded to port 1 in VLAN 1 and ports 9, 10, and 11 in VLAN 5.

Each VLAN maintains its own logical address table. A known address in one VLAN is unknown in another until it is learned or statically associated with a port in the other VLAN. An address can be secure in one VLAN and dynamic in another. An address can be static in one VLAN and dynamic in another.

Changing the Address Aging Time

Dynamic addresses are source MAC addresses that the switch learns and then drops when they are not in use. The aging time parameter defines how long the switch retains unseen addresses in the table. This parameter applies to all VLANs.

Setting too short an aging time can cause addresses to be prematurely removed from the table. Then when the switch receives a packet for an unknown destination, it floods the packet to all ports in the same VLAN as the receiving port. This unnecessary flooding can impact performance. Setting too long an aging time can cause the address table to be filled with unused addresses; it can cause delays in establishing connectivity when a workstation is moved to a new port.

Beginning in privileged EXEC mode, follow these steps to configure the dynamic address table aging time.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac-address-table aging-time <i>seconds</i>	Enter the number of seconds that dynamic addresses are to be retained in the address table. You can enter a number from 10 to 1000000.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac-address-table aging-time	Verify your entry.

Removing Dynamic Address Entries

Beginning in privileged EXEC mode, follow these steps to remove a dynamic address entries:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clear mac-address-table dynamic [address <i>mac-addr</i> interface <i>interface-id</i> vlan <i>vlan-id</i>]	(Optional) Enter the address <i>mac-addr</i> to delete the specified MAC address. (Optional) Enter the interface <i>interface-id</i> to delete all dynamic MAC addresses on the specified physical port or port channel. (Optional) Enter the vlan <i>vlan-id</i> to delete all dynamic MAC addresses for the specified VLAN. Valid IDs are from 1 to 1005; do not enter leading zeroes.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac-address-table	Verify your entry.

MAC Address Notification

MAC address notification enables you to keep track of the MAC addresses that are learned or removed from your switch.

When a new MAC address is learned or an old MAC address is removed from the switch, an SNMP notification (trap) is generated. Traps can be bundled and sent at regular intervals.

Enabling Notification of Learned or Deleted MAC Addresses

You can enable the MAC notification feature on the switch. The MAC notification feature can bundle SNMP traps and send them to the CMS at regular intervals.

Beginning in privileged EXEC mode, follow these steps to enable the MAC address notification feature:

	Command	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>mac-address-table notification</code> <code>[interval seconds historysize value]</code>	Enable the MAC address notification feature. For interval seconds, the range is 0 to 2147483647. The default is 1 second. The switch sends the notification trap after the interval setting has expired. For history size, the range is 0 to 500 entries. The default is 1 entry.
Step 3	<code>SNMP-server enable traps</code> <code>mac-notification</code>	Enable SNMP notification of MAC address additions and deletions.
Step 4	<code>interface interface-id</code>	Enter interface configuration mode for the port that you want to configure.
Step 5	<code>SNMP trap mac-notification</code> <code>[added removed]</code>	Enable or disable MAC address traps on the port.
Step 6	<code>end</code>	Return to privileged EXEC mode.
Step 7	<code>show mac-address notification</code> <code>table notification interface</code> <code>interface-id</code>	Verify your settings.

To disable the switch from sending MAC address notification traps, use the **no snmp-server enable traps mac-notification** global configuration command. To disable the MAC address notification traps on a specific interface, use the **no snmp trap mac-notification** interface configuration command. To disable the MAC address notification feature, use the **no mac-address-table notification** global configuration command.

This example shows how to specify 172.20.10.10 as the NMS, enable the switch to send MAC address notification traps to the NMS, enable the MAC address notification feature, set the interval time to 60 seconds, set the history-size to 100 entries, and enable traps whenever a MAC address is added on Fast Ethernet interface 0/4.

```
Switch(config)# snmp-server host 172.20.10.10
Switch(config)# snmp-server enable traps mac-notification
Switch(config)# mac-address-table notification interval 60 history-size 100
Switch(config)# interface fastethernet0/4
Switch(config-if)# snmp trap mac-notification added
```

You can verify the previous commands by entering the **show mac-address-table notification** privileged EXEC command.

Adding Secure Addresses

The secure address table contains secure MAC addresses and their associated ports and VLANs. A secure address is a manually entered unicast address or dynamically learnt address that is forwarded to only one port per VLAN. If you enter a static address that is already assigned to another port, the request will be rejected.

Secure addresses can be learned dynamically if the configured secure addresses do not reach the maximum limit of the port.

Beginning in privileged EXEC mode, follow these steps to add a secure address:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Identify a specific interface for configuration, and enter interface configuration mode.
Step 3	switchport port-security mac address <i>mac-address</i>	Add a secure address.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mac-address-table secure	Verify your entry.

Removing Secure Addresses

Beginning in privileged EXEC mode, follow these steps to remove a secure address:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no switchport port-security mac address <i>mac-address</i>	Remove a secure address.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac-address-table secure	Verify your entry.

You can remove all secure addresses by using the **clear mac-address-table secure** command in privileged EXEC mode.

Adding and Removing Static Address Entries

A static address has these characteristics:

- It is manually entered in the address table and must be manually removed.
- It can be a unicast or multicast address.
- It does not age and is retained when the switch restarts.

You can add and remove static addresses and define the forwarding behavior for them. The forwarding behavior determines how a port that receives a packet forwards it to another port for transmission. Because all ports are associated with at least one VLAN, the switch acquires the VLAN ID for the address from the ports that you specify. You can specify a different list of destination ports for each source port.

A static address in one VLAN must be a static address in other VLANs. A packet with a static address that arrives on a VLAN where it has not been statically entered is flooded to all ports and not learned.

You add a static address to the address table by specifying the destination MAC address (unicast or multicast) and the VLAN from which it is received. Packets received with this destination address are forwarded to the interface specified with the interface-id option.

Beginning in privileged EXEC mode, follow these steps to add a static address:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac-address-table static <i>mac-addr</i> vlan <i>vlan-id</i> interface <i>interface-id</i> <i>interface-id</i> ...	Add a static address to the mac address table: <ul style="list-style-type: none"> • For <i>mac-addr</i>, specify the destination MAC address (unicast or multicast) to add to the address table. Packets with this destination address received in the specified VLAN are forwarded to the specified interface. • For <i>vlan-id</i>, specify the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are 1 to 1005; do not enter leading zeroes. • For <i>interface-id...</i>, specify the interface to which the received packet is forwarded. Valid interfaces include physical ports and EtherChannel port-channels. Multiple interfaces can be specified for multicast addresses.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac-address-table static	Verify your entry.

To remove static entries from the address table, use the **no mac-address-table static** *mac-addr* **vlan** *vlan-id* **interface** *interface-id* *interface-id...* global configuration command.

Configuring Static Addresses for EtherChannel Port Groups

Follow these rules if you are configuring a static address to forward to ports in an EtherChannel port group:

- For default source-based port groups, configure the static address to forward to *all ports* in the port group to eliminate lost packets.
- For destination-based port groups, configure the address to forward to *only one port* in the port group to avoid the transmission of duplicate packets.

Configuring TACACS+

You can use the Terminal Access Controller Access Control System Plus (TACACS+) to manage network security (authentication, authorization, and accounting [AAA]) from a server. This section describes how TACACS+ works and how you can configure it. For complete syntax and usage information for the commands described in this chapter, refer to the *Cisco IOS Release 12.1 Security Command Reference*.

You can only configure this feature by using the CLI; you cannot configure it through the Cluster Management Suite.

**Note**

If TACACS+ is configured on the command switch, TACACS+ must also be configured on all member switches to access the switch cluster from CMS. For more information about switch clusters, see [Chapter 5, “Clustering Switches.”](#)

In large enterprise networks, the task of administering passwords on each device can be simplified by centralizing user authentication on a server. TACACS+ is an access-control protocol that allows a switch to authenticate all login attempts through a central server. The network administrator configures the switch with the address of the TACACS+ server, and the switch and the server exchange messages to authenticate each user before allowing access to the management console.

TACACS+ consists of three services: authentication, authorization, and accounting. Authentication determines who the user is and whether or not the user is allowed access to the switch. Authorization determines what the user is allowed to do on the system. Accounting collects data related to resource usage.

The TACACS+ feature is disabled by default. However, you can enable and configure it by using the CLI. You can access the CLI through the console port or through Telnet. To prevent a lapse in security, you cannot configure TACACS+ through a network-management application. When enabled, TACACS+ can authenticate users accessing the switch through the CLI.

**Note**

Although the TACACS+ configuration is performed through the CLI, the TACACS+ server authenticates HTTP connections that have been configured with a privilege level of 15.

Configuring the TACACS+ Server Host

Use the **tacacs-server host** privileged EXEC command to specify the names of the IP host or hosts maintaining an AAA/TACACS+ server. On TACACS+ servers, you can configure these additional options:

- Number of seconds that the switch waits while trying to contact the server before timing out.
- Encryption key to encrypt and decrypt all traffic between the router and the daemon.
- Number of attempts that a user can make when entering a command that is being authenticated by TACACS+.

Beginning in privileged EXEC mode, follow these steps to configure the TACACS+ server.

	Command	Purpose
Step 1	tacacs-server host <i>name</i> [timeout <i>integer</i>] [key <i>string</i>]	Define a TACACS+ host. Entering the timeout and key parameters with this command overrides the global values that you can enter with the tacacs-server timeout (Step 3) and the tacacs-server key commands (Step 5).
Step 2	tacacs-server retransmit <i>retries</i>	Enter the number of times the server searches the list of TACACS+ servers before stopping. The default is two.
Step 3	tacacs-server timeout <i>seconds</i>	Set the interval that the server waits for a TACACS+ server host to reply. The default is 5 seconds.
Step 4	tacacs-server attempts <i>count</i>	Set the number of login attempts that can be made on the line.
Step 5	tacacs-server key <i>key</i>	Define a set of encryption keys for all of TACACS+ and communication between the access server and the TACACS daemon. Repeat the command for each encryption key.
Step 6	exit	Return to privileged EXEC mode.
Step 7	show tacacs	Verify your entries.

Configuring Login Authentication

Beginning in privileged EXEC mode, follow these steps to configure login authentication by using AAA/TACACS+:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA/TACACS+.
Step 3	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2...</i>]	Enable authentication at login, and create one or more lists of authentication methods.
Step 4	line [aux console tty vty] <i>line-number</i> [<i>ending-line-number</i>]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.
Step 5	login authentication { default <i>list-name</i> }	Apply the authentication list to a line or set of lines.
Step 6	exit	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.

The variable *list-name* is any character string used to name the list you are creating. The *method* variable refers to the actual methods the authentication algorithm tries, in the sequence entered. You can choose one of these methods:

- **line**—Uses the line password for authentication. You must define a line password before you can use this authentication method. Use the **password** *password* line configuration command.
- **local**—Uses the local username database for authentication. You must enter username information in the database. Use the **username** *password* global configuration command.

- **tacacs+**—Uses TACACS+ authentication. You must configure the TACACS+ server before you can use this authentication method. For more information, see the “Configuring the TACACS+ Server Host” section on page 6-20.

To create a default list that is used if **no list** is specified in the **login authentication** line configuration command, use the **default** keyword followed by the methods that you want used in default situations.

The additional methods of authentication are used only if previous method returns an error, not if it fails. To specify that the authentication should succeed even if all methods return an error, specify **none** as the final method in the command line.

Specifying TACACS+ Authorization for Privileged EXEC Access and Network Services

You can use the **aaa authorization** global configuration command with the **tacacs+** keyword to set parameters that restrict a user’s network access to Cisco IOS privileged-mode (EXEC access) and to network services such as Serial Line Internet Protocol (SLIP), Point-to-Point Protocol (PPP) with Network Control Protocols (NCPs), and AppleTalk Remote Access (ARA).

The **aaa authorization exec tacacs+ local** command sets these authorization parameters:

- Uses TACACS+ for privileged EXEC access authorization if authentication was done by using TACACS+.
- Uses the local database if authentication was not done by using TACACS+.



Note

Authorization is bypassed for authenticated users who login through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify TACACS+ authorization for privileged EXEC access and network services:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa authorization network tacacs+	Configure the switch for user TACACS+ authorization for all network-related service requests, including SLIP, PPP NCPs, and ARA protocols.
Step 3	aaa authorization exec tacacs+	Configure the switch for user TACACS+ authorization to determine if the user is allowed privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
Step 4	exit	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.

Starting TACACS+ Accounting

You use the **aaa accounting** command with the **tacacs+** keyword to turn on TACACS+ accounting for each Cisco IOS privilege level and for network services.

Beginning in privileged EXEC mode, follow these steps to enable TACACS+ accounting:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa accounting exec start-stop tacacs+	Enable TACACS+ accounting to send a start-record accounting notice at the beginning of an EXEC process and a stop-record at the end.
Step 3	aaa accounting network start-stop tacacs+	Enable TACACS+ accounting for all network-related service requests, including SLIP, PPP, and PPP NCPs.
Step 4	exit	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.

**Note**

These commands are documented in the “Accounting and Billing Commands” chapter of the *Cisco IOS Release 12.1 Security Command Reference*.

Configuring a Switch for Local AAA

You can configure AAA to operate without a server by setting the switch to implement AAA in local mode. The switch then verifies authentication and authorization. No accounting is available in this configuration.

Beginning in privileged EXEC mode, follow these steps to configure the switch for local AAA:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication login default local	Set the login authorization to default to local.
Step 4	aaa authorization exec local	Configure user AAA authorization for all network-related service requests, including SLIP, PPP NCPs, and ARA protocols.
Step 5	aaa authorization network local	Configure user AAA authorization to determine if the user is allowed to run a privileged EXEC shell.
Step 6	username name privilege level password password	Enter the local database. Repeat this command for each user.
Step 7	show running-config	Verify your entries.

Controlling Switch Access with RADIUS

This section describes how to enable and configure the Remote Authentication Dial-In User Service (RADIUS), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. RADIUS is facilitated through authentication, authorization, accounting (AAA) and can be enabled only through AAA commands.

**Note**

For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Release 12.1*.

This section contains this configuration information:

- [Understanding RADIUS, page 6-24](#)
- [RADIUS Operation, page 6-25](#)
- [Configuring RADIUS, page 6-26](#)
- [Displaying the RADIUS Configuration, page 6-37](#)

Understanding RADIUS

RADIUS is a distributed client/server system that secures networks against unauthorized access. RADIUS clients run on supported Cisco routers and switches (including Catalyst 3550 multilayer switches and Catalyst 2950 switches) and send authentication requests to a central RADIUS server, which contains all user authentication and network service access information. The RADIUS host is normally a multiuser system running RADIUS server software from Cisco (Cisco Secure Access Control Server version 3.0), Livingston, Merit, Microsoft, or another software provider. For more information, refer to the RADIUS server documentation.


Use RADIUS in these network environments that require access security:

- Networks with multiple-vendor access servers, each supporting RADIUS. For example, access servers from several vendors use a single RADIUS server-based security database. In an IP-based network with multiple vendors' access servers, dial-in users are authenticated through a RADIUS server that has been customized to work with the Kerberos security system.
- Turnkey network security environments in which applications support the RADIUS protocol, such as in an access environment that uses a *smart card* access control system. In one case, RADIUS has been used with Enigma's security cards to validate users and to grant access to network resources.
- Networks already using RADIUS. You can add a Cisco switch containing a RADIUS client to the network. This might be the first step when you make a transition to a TACACS+ server.
- Network in which the user must only access a single service. Using RADIUS, you can control user access to a single host, to a single utility such as Telnet, or to the network through a protocol such as IEEE 802.1X. For more information about this protocol, see [Chapter 7, "Configuring 802.1X Port-Based Authentication."](#)
- Networks that require resource accounting. You can use RADIUS accounting independently of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, showing the amount of resources (such as time, packets, bytes, and so forth) used during the session. An Internet service provider might use a freeware-based version of RADIUS access control and accounting software to meet special security and billing needs.

RADIUS is not suitable in these network security situations:

- Multiprotocol access environments. RADIUS does not support AppleTalk Remote Access (ARA), NetBIOS Frame Control Protocol (NBFCP), NetWare Asynchronous Services Interface (NASI), or X.25 PAD connections.
- Switch-to-switch or router-to-router situations. RADIUS does not provide two-way authentication. RADIUS can be used to authenticate from one device to a non-Cisco device if the non-Cisco device requires authentication.
- Networks using a variety of services. RADIUS generally binds a user to one service model.

Figure 6-5 Typical AAA Network Configuration



RADIUS Operation

When a user attempts to log in and authenticate to a switch that is access controlled by a RADIUS server, these events occur:

1. The user is prompted to enter a username and password.
2. The username and encrypted password are sent over the network to the RADIUS server.
3. The user receives one of these responses from the RADIUS server:
 - a. ACCEPT—The user is authenticated.
 - b. REJECT—The user is either not authenticated and is prompted to re-enter the username and password, or access is denied.
 - c. CHALLENGE—A challenge requires additional data from the user.
 - d. CHALLENGE PASSWORD—A response requests the user to select a new password.

The ACCEPT or REJECT response is bundled with additional data that is used for privileged EXEC or network authorization. Users must first successfully complete RADIUS authentication before proceeding to RADIUS authorization, if it is enabled. The additional data included with the ACCEPT or REJECT packets includes these items:

- Telnet, rlogin, or privileged EXEC services
- Connection parameters, including the host or client IP address, access list, and user timeouts

Configuring RADIUS

This section describes how to configure your switch to support RADIUS. At a minimum, you must identify the host or hosts that run the RADIUS server software and define the method lists for RADIUS authentication. You can optionally define method lists for RADIUS authorization and accounting.

A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used (such as TACACS+ or local username lookup), thus ensuring a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

You should have access to and should configure a RADIUS server before configuring RADIUS features on your switch.

This section contains this configuration information:

- [Default RADIUS Configuration, page 6-26](#)
- [Identifying the RADIUS Server Host, page 6-27](#) (required)
- [Configuring RADIUS Login Authentication, page 6-29](#) (required)
- [Defining AAA Server Groups, page 6-31](#) (optional)
- [Configuring RADIUS Authorization for Privileged EXEC Access and Network Services, page 6-33](#) (optional)
- [Starting RADIUS Accounting, page 6-34](#) (optional)
- [Configuring Settings for All RADIUS Servers, page 6-35](#) (optional)
- [Configuring the Switch to Use Vendor-Specific RADIUS Attributes, page 6-35](#) (optional)
- [Configuring the Switch for Vendor-Proprietary RADIUS Server Communication, page 6-36](#) (optional)

Default RADIUS Configuration

RADIUS and AAA are disabled by default.

To prevent a lapse in security, you cannot configure RADIUS through a network management application. When enabled, RADIUS can authenticate users accessing the switch through the CLI.

Identifying the RADIUS Server Host

Switch-to-RADIUS-server communication involves several components:

- Host name or IP address
- Authentication destination port
- Accounting destination port
- Key string
- Timeout period
- Retransmission value

You identify RADIUS security servers by their host name or IP address, host name and specific UDP port numbers, or their IP address and specific UDP port numbers. The combination of the IP address and the UDP port number creates a unique identifier, allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. This unique identifier enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address.

If two different host entries on the same RADIUS server are configured for the same service—for example, accounting—the second host entry configured acts as a fail-over backup to the first one. Using this example, if the first host entry fails to provide accounting services, the switch tries the second host entry configured on the same device for accounting services. (The RADIUS host entries are tried in the order that they are configured.)

A RADIUS server and the switch use a shared secret text string to encrypt passwords and exchange responses. To configure RADIUS to use the AAA security commands, you must specify the host running the RADIUS server daemon and a secret text (key) string that it shares with the switch.

The timeout, retransmission, and encryption key values can be configured globally for all RADIUS servers, on a per-server basis, or in some combination of global and per-server settings. To apply these settings globally to all RADIUS servers communicating with the switch, use the three unique global configuration commands: **radius-server timeout**, **radius-server retransmit**, and **radius-server key**. To apply these values on a specific RADIUS server, use the **radius-server host** global configuration command.



Note

If you configure both global and per-server functions (timeout, retransmission, and key commands) on the switch, the per-server timer, retransmission, and key value commands override global timer, retransmission, and key value commands. For information on configuring these setting on all RADIUS servers, see the [“Configuring Settings for All RADIUS Servers”](#) section on page 6-35.

You can configure the switch to use AAA server groups to group existing server hosts for authentication. For more information, see the [“Defining AAA Server Groups”](#) section on page 6-31.

Beginning in privileged EXEC mode, follow these steps to configure per-server RADIUS server communication. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } [auth-port <i>port-number</i>] [acct-port <i>port-number</i>] [timeout <i>seconds</i>] [retransmit <i>retries</i>] [key <i>string</i>]	<p>Specify the IP address or host name of the remote RADIUS server host.</p> <ul style="list-style-type: none"> (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before retransmitting. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set with the radius-server host, the setting of the radius-server timeout global configuration command is used. (Optional) For retransmit retries, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host command, the setting of the radius-server retransmit global configuration command is used. (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified RADIUS server, use the **no radius-server host** *hostname* | *ip-address* global configuration command.

This example shows how to configure one RADIUS server to be used for authentication and another to be used for accounting:

```
Switch(config)# radius-server host 172.29.36.49 auth-port 1612 key rad1
Switch(config)# radius-server host 172.20.36.50 acct-port 1618 key rad2
```

This example shows how to configure *host1* as the RADIUS server and to use the default ports for both authentication and accounting:

```
Switch(config)# radius-server host host1
```

**Note**

You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch. For more information, refer to the RADIUS server documentation.

Configuring RADIUS Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific interface before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure login authentication. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2...</i>]	<p>Create a login authentication method list.</p> <ul style="list-style-type: none"> To create a default list that is used when a named list is <i>not</i> specified in the login authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. For <i>list-name</i>, specify a character string to name the list you are creating. For <i>method1...</i>, specify the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. <p>Select one of these methods:</p> <ul style="list-style-type: none"> line—Use the line password for authentication. You must define a line password before you can use this authentication method. Use the password <i>password</i> line configuration command. local—Use the local username database for authentication. You must enter username information in the database. Use the username <i>password</i> global configuration command. radius—Use RADIUS authentication. You must configure the RADIUS server before you can use this authentication method. For more information, see the “Identifying the RADIUS Server Host” section on page 6-27.
Step 4	line [console tty vty] <i>line-number</i> [<i>ending-line-number</i>]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.
Step 5	login authentication { default <i>list-name</i> }	<p>Apply the authentication list to a line or set of lines.</p> <ul style="list-style-type: none"> If you specify default, use the default list created with the aaa authentication login command. For <i>list-name</i>, specify the list created with the aaa authentication login command.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable AAA authentication, use the **no aaa authentication login** { **default** | *list-name* } *method1* [*method2...*] global configuration command. To either disable RADIUS authentication for logins or to return to the default value, use the **no login authentication** { **default** | *list-name* } line configuration command.

Defining AAA Server Groups

You can configure the switch to use AAA server groups to group existing server hosts for authentication. You select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list, which lists the IP addresses of the selected server hosts.

Server groups also can include multiple host entries for the same server if each entry has a unique identifier (the combination of the IP address and UDP port number), allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. If you configure two different host entries on the same RADIUS server for the same service, (for example, accounting), the second configured host entry acts as a fail-over backup to the first one.

You use the **server** group server configuration command to associate a particular server with a defined group server. You can either identify the server by its IP address or identify multiple host instances or entries by using the optional **auth-port** and **acct-port** keywords.

Beginning in privileged EXEC mode, follow these steps to define the AAA server group and associate a particular RADIUS server with it:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } [auth-port <i>port-number</i>] [acct-port <i>port-number</i>] [timeout <i>seconds</i>] [retransmit <i>retries</i>] [key <i>string</i>]	<p>Specify the IP address or host name of the remote RADIUS server host.</p> <ul style="list-style-type: none"> (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before retransmitting. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set, with the radius-server host global configuration command, the setting of the radius-server timeout global configuration command is used. (Optional) For retransmit <i>retries</i>, specify the number of times a RADIUS request is resent to a server, if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host global configuration command, the setting of the radius-server retransmit global configuration command is used. (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
Step 3	aaa new-model	Enable AAA.
Step 4	aaa group server radius <i>group-name</i>	Define the AAA server-group with a group name. This command puts the switch in a server group configuration mode.
Step 5	server <i>ip-address</i>	Associate a particular RADIUS server with the defined server group. Repeat this step for each RADIUS server in the AAA server group. Each server in the group must be previously defined in Step 2.
Step 6	end	Return to privileged EXEC mode.

	Command	Purpose
Step 7	<code>show running-config</code>	Verify your entries.
Step 8	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.
Step 9		Enable RADIUS login authentication. See the “ Configuring RADIUS Login Authentication ” section on page 6-29.

To remove the specified RADIUS server, use the **no radius-server host** *hostname | ip-address* global configuration command. To remove a server group from the configuration list, use the **no aaa group server radius** *group-name* global configuration command. To remove the IP address of a RADIUS server, use the **no server ip-address** server group configuration command.

In this example, the switch is configured to recognize two different RADIUS group servers (*group1* and *group2*). Group1 has two different host entries on the same RADIUS server configured for the same services. The second host entry acts as a fail-over backup to the first entry.

```
Switch(config)# radius-server host 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config)# radius-server host 172.10.0.1 auth-port 1645 acct-port 1646
Switch(config)# aaa new-model
Switch(config)# aaa group server radius group1
Switch(config-sg-radius)# server 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config-sg-radius)# exit
Switch(config)# aaa group server radius group2
Switch(config-sg-radius)# server 172.20.0.1 auth-port 2000 acct-port 2001
Switch(config-sg-radius)# exit
```

Configuring RADIUS Authorization for Privileged EXEC Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user’s profile, which is in either the local user database or on the security server, to configure the user’s session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **radius** keyword to set parameters that restrict a user’s network access to privileged EXEC mode.

The **aaa authorization exec radius local** command sets these authorization parameters:

- Use RADIUS for privileged EXEC access authorization if authentication was performed by using RADIUS.
- Use the local database if authentication was not performed by using RADIUS.



Note

Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify RADIUS authorization for privileged EXEC access and network services:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa authorization network radius	Configure the switch for user RADIUS authorization for all network-related service requests.
Step 3	aaa authorization exec radius	Configure the switch for user RADIUS authorization to determine if the user has privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable authorization, use the **no aaa authorization {network | exec} method1** global configuration command.

Starting RADIUS Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Beginning in privileged EXEC mode, follow these steps to enable RADIUS accounting for each Cisco IOS privilege level and for network services:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa accounting network start-stop radius	Enable RADIUS accounting for all network-related service requests.
Step 3	aaa accounting exec start-stop radius	Enable RADIUS accounting to send a start-record accounting notice at the beginning of an privileged EXEC process and a stop-record at the end.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable accounting, use the **no aaa accounting {network | exec} {start-stop} method1...** global configuration command.

Configuring Settings for All RADIUS Servers

Beginning in privileged EXEC mode, follow these steps to configure global communication settings between the switch and all RADIUS servers:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server key <i>string</i>	Specify the shared secret text string used between the switch and all RADIUS servers. Note The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 3	radius-server retransmit <i>retries</i>	Specify the number of times the switch sends each RADIUS request to the server before giving up. The default is 3; the range 1 to 1000.
Step 4	radius-server timeout <i>seconds</i>	Specify the number of seconds a switch waits for a reply to a RADIUS request before sending the request. The default is 5 seconds; the range is 1 to 1000.
Step 5	radius-server deadtime <i>minutes</i>	Specify the number of minutes a RADIUS server, which is not responding to authentication requests, to be skipped, thus avoiding the wait for the request to timeout before trying the next configured server. The default is 0; the range is 1 to 1440 minutes.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your settings.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting for the retransmit, timeout, and deadtime, use the **no** forms of these commands.

Configuring the Switch to Use Vendor-Specific RADIUS Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific information between the switch and the RADIUS server by using the vendor-specific attribute (attribute 26). Vendor-specific attributes (VSAs) allow vendors to support their own extended attributes not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option by using the format recommended in the specification. Cisco's vendor-ID is 9, and the supported option has vendor-type 1, which is named *cisco-avpair*. The value is a string with this format:

```
protocol : attribute sep value *
```

Protocol is a value of the Cisco *protocol* attribute for a particular type of authorization. *Attribute* and *value* are an appropriate attribute-value (AV) pair defined in the Cisco TACACS+ specification, and *sep* is = for mandatory attributes and * for optional attributes. This allows the full set of features available for TACACS+ authorization to also be used for RADIUS.

For example, this AV pair activates Cisco's *multiple named ip address pools* feature during IP authorization (during PPP's IPCP address assignment):

```
cisco-avpair= "ip:addr-pool=first"
```

This example shows how to provide a user logging in from a switch with immediate access to privileged EXEC commands:

```
cisco-avpair= "shell:priv-lvl=15"
```

Other vendors have their own unique vendor-IDs, options, and associated VSAs. For more information about vendor-IDs and VSAs, refer to RFC 2138, "Remote Authentication Dial-In User Service (RADIUS)."

Beginning in privileged EXEC mode, follow these steps to configure the switch to recognize and use VSAs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server vsa send [accounting authentication]	<p>Enable the switch to recognize and use VSAs as defined by RADIUS IETF attribute 26.</p> <ul style="list-style-type: none"> • (Optional) Use the accounting keyword to limit the set of recognized vendor-specific attributes to only accounting attributes. • (Optional) Use the authentication keyword to limit the set of recognized vendor-specific attributes to only authentication attributes. <p>If you enter this command without keywords, both accounting and authentication vendor-specific attributes are used.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your settings.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For a complete list of RADIUS attributes or more information about vendor-specific attribute 26, refer to the "RADIUS Attributes" appendix in the *Cisco IOS Security Configuration Guide for Release 12.1*.

Configuring the Switch for Vendor-Proprietary RADIUS Server Communication

Although an IETF draft standard for RADIUS specifies a method for communicating vendor-proprietary information between the switch and the RADIUS server, some vendors have extended the RADIUS attribute set in a unique way. Cisco IOS software supports a subset of vendor-proprietary RADIUS attributes.

As mentioned earlier, to configure RADIUS (whether vendor-proprietary or IETF draft-compliant), you must specify the host running the RADIUS server daemon and the secret text string it shares with the switch. You specify the RADIUS host and secret text string by using the **radius-server** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to specify a vendor-proprietary RADIUS server host and a shared secret text string:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } non-standard	Specify the IP address or host name of the remote RADIUS server host and identify that it is using a vendor-proprietary implementation of RADIUS.
Step 3	radius-server key <i>string</i>	Specify the shared secret text string used between the switch and the vendor-proprietary RADIUS server. The switch and the RADIUS server use this text string to encrypt passwords and exchange responses. Note The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your settings.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the vendor-proprietary RADIUS host, use the **no radius-server host** {*hostname* | *ip-address*} **non-standard** command. To disable the key, use the **no radius-server key** command.

This example shows how to specify a vendor-proprietary RADIUS host and to use a secret key of *rad124* between the switch and the server:

```
Switch(config)# radius-server host 172.20.30.15 nonstandard
Switch(config)# radius-server key rad124
```

Displaying the RADIUS Configuration

To display the RADIUS configuration, use the **show running-config** privileged EXEC command.

